# High-Frequency Exchange Rate Forecasting

Charlie X. Cai [a]
University of Leeds

Qi Zhang * [b]
University of Leeds

**Abstract**

Predictability of exchange rate movement is of great interest to both practitioners and regulators. We examine the predictability of exchange rate movement in the high-frequency domain. To this end, we apply a model designed for modelling high-frequency and irregularly spaced data, the autoregressive conditional multinomial–autoregressive conditional duration (ACM–ACD) model. Studying three pairs of currencies, we find strong predictability in the high-frequency quote change data, with the rate of correct predictions varying from 54 to 70%. We demonstrate that filtering the data, by increasing the threshold of mid-quote price change, in combination with dynamic learning, can improve forecasting performance.

JEL classification: F31, F37, G17

*Keywords*: foreign exchange, high-frequency data, forecasting, duration model.

[a] *Professor of Finance, Leeds University Business School, LS2 9JT, Leeds, UK*

[b] *Associate Professor in Accounting and Finance, Leeds University Business School, LS2 9JT, Leeds, UK.*

* We thank John Doukas (the editor) and an anonymous referee for their constructive and very helpful comments. We also thank Matthias Böhm the discussion and the participants at the 2nd International conference of the Financial Engineering and Banking Society for their helpful comments.

Corresponding Author: Qi Zhang, q.zhang@leeds.ac.uk.

# 1 Introduction

Foreign exchange rate forecasting is an ongoing challenge to academics, practitioners and policy makers. A large number of studies have been dedicated to the examination of exchange rate dynamics. These studies have produced at least two main conclusions among other insights. First, although macroeconomic fundamentals should in theory dictate the purchasing power of a country's currency, they are less useful in predicting exchange rates, given the nature of low-frequency announcements.[1] Second, better information sets or modelling techniques may help improve short-term forecasting of exchange rate movements; for example, the use of customer order-flow information.[2] While the random walk model is difficult to surpass in forecasting the mean, Hong *et al.* (2007) show that sophisticated time-series models are preferable for out-of-sample density forecasting; however, there are some difficulties in the direct application of such models.[3] In this regard, a successful model for forecasting the mean is much more useful practically, since it can be directly applied to constructing quotation strategy for dealers and trading strategy for traders.

To meet the challenge of finding a suitable model for foreign exchange (FX) rate forecasting, we explore the use of the autoregressive conditional multinomial–autoregressive conditional duration (ACM–ACD) model in high-frequency FX rate forecasting, as introduced by Russell and Engle (2005). We consider this model to be a potentially useful tool for high-frequency exchange rate forecasting for three reasons.

---

[1] See for example, Sarno and Taylor (2002) or Taylor and Taylor (2004) for surveys on the purchasing power parity debate. It is also widely documented that traditional exchange rate determination models perform poorly in explaining and forecasting exchange rates (Meese and Rogoff, 1983; Cheung *et al.*, 2005; Engel *et al.*, 2008).

[2] Motivated by the market microstructure literature, Evans and Lyons (2002, 2005) show empirically that order flow is an important determinant of daily exchange rates. Several studies in this area try to bring macroeconomic and microstructure approaches together to study exchange rates (Bacchetta and van Wincoop, 2004; Evans and Lyons, 2007; Rime *et al.*, 2010). These studies show that the short-term determinants of exchange rate movements are more likely driven by the short-term supply and demand fluctuations of the exchange rate market.

[3] Sarno and Valente (2005) were the first to document that complex (non-linear) models of exchange rates produce superior density forecasts of exchange rates. Hong *et al.*, (2007) suggest that a better density forecast can be used in many ways (e.g. risk management, option pricing). However, they also point out that direct application of these models and evaluation of their performance requires further exploration.

First, following the technological advances of the last two decades, high-frequency trading at the intra-day level has become popular and technical trading rules embedded in algorithmic trading play an important role in financial markets. Indeed, high-frequency algorithmic trading has been judged one of the major drivers behind the growth of FX market turnover in the latest Bank for International Settlements (BIS, 2013) report on global FX market activity. However, most of the previous literature on intra-day exchange rate forecasting has focused on regular time intervals such as 30 minutes or one hour. None of the existing studies consider tick-by-tick frequency with irregular time spacing.[4] Our study fills this void by modelling the exchange rate movement tick by tick and evaluating the model's out-of-sample forecasting performance.

Second, there is empirical evidence to suggest that the ACM–ACD model performs well in high-frequency forecasting. Zhang *et al.* (2009) demonstrate, in the context of the US equity market, that it provides better forecasting performance than either the asymmetric ACD model (AACD) – an alternative asymmetric extension of the ACD model proposed by Bauwens and Giot (2003) – or the random walk model. They also show that the forecast performance of the model generally decreases as the liquidity of the stock increases, with the exception of the most liquid stocks. This suggests that there is a potential gain in forecasting performance by applying this model to the exchange rate market, the most liquid of financial markets.

Finally, the ACM–ACD model is designed for modelling irregularly spaced high-frequency dynamics. It incorporates an important set of information, the durations between price changes, into modelling price dynamics. Theoretical studies (e.g., Easley and O'Hara, 1987) in market microstructure show that the time between intra-day market events, like trades and quote updates, may contain information about market states. Dealers receive information from their customer orders and they also learn about market-wide order flow from brokered inter-dealer trades (Loyns, 2001: Chapter 3). Given that order-flow information is valuable and private

[4] One exception is Engle and Russell's (1997) study, which examines high-frequency quote changes in the foreign exchange rate market using an ACD model. However, their study does not produce quote change forecasts, since it models only the duration and not the direction of quote updates.

to the dealers, in this study we explore whether a sophisticated model utilizing both duration and direction of quote change information, both of which are publicly available, would help to abstract private information and produce better out-of-sample forecasting of quote updates.[5]

Specifically, we examine the following two research questions. (1) Does the ACM–ACD model provide better forecasting of the high-frequency exchange rate than a random walk model? (2) Can the forecasting provided by the ACM–ACD model be put to practical use by dealers and traders?

We examine three currencies against the dollar (USD) for the year 2010: the euro (EUR), the British pound (GBP) and the Japanese yen (JPY). Using high-frequency bid and ask quotes for spot FX rates and daily/weekly estimation windows with a daily rolling frequency, we assess out-of-sample forecasting by the ACM–ACD model at the intra-day tick-by-tick level.

Our analysis yields the following key findings. First, the ACM–ACD model provides much improved forecasting over that of the benchmark random walk model and others reported in the existing literature. The forecast accuracy of the models ranges from 55 to 63% for the three currency pairs when no data filter is applied. The correlation analyses, Wald test and Clark–West (2006) test of the relationship between the predicted and realized quote movements also confirm the positive and significant predictive power of these models against the random walk benchmark. We further show that when a modest filter is applied to the raw data (by increasing the tick size threshold for counting the number of quote changes), the prediction accuracy increases to between 59 and 70%. However, when a higher threshold of price movement is imposed, the prediction accuracy drops to a range between 53 and 54%. This suggests that over-filtering may reduce the useful information within the data. The conclusions drawn from the filtered results of the full sample period are potentially biased towards hindsight.[6] To avoid look-back bias in the choice of threshold and to include an element of learning, we construct a dynamic switching strategy updated by selecting the best individual strategy, based on past

[5] A similar idea has been adopted in the intra-day risk management literature. For example, Dionne *et al.* (2009) propose an Intraday Value at Risk measurement by modelling the joint density of trade durations and high-frequency returns, which shows reliable performance of measuring intra-day risk.

[6] We thank the referee for pointing out this issue and making helpful suggestions.

performance of the available filtering strategies. Our results show that, by dealing with the uncertainty of filtering choice in this way, while the best-performing filtering strategies may not be achievable ex ante, the dynamic strategy still improves performance over the non-filtered models.

Second, we find that trading strategies based on the directional forecasts can earn substantial returns before transaction costs. Specifically, we simulate a portfolio switching the position from longing one currency to longing the other when the directional forecasts change, while holding the position when same-signed forecasts are produced. This strategy can generate a strikingly large, if somewhat theoretical, average daily return, accumulated from intra-day transactions, ranging from 5 to 15%, from the prediction models before costs. These returns are much better than the buy-and-hold return, which is approximately 0.03% daily. While applying filters improves directional prediction accuracy, it does not always produce greater economic benefit. Only the return in the GBP/USD series is improved after applying data filters. Our risk-adjustment tests based on the certainty-equivalent return produce similar conclusions to the portfolio return measures.

Third, we document the overwhelming impact of transaction costs on trading strategy. The documented before-cost return requires high-frequency automatic trading; specifically, of about eight trades per minute, according to our findings. When transaction costs of half a basis point are applied to each intra-day transaction, the profit quickly turns to loss for most of the strategies, except for the two JPY models using filtered data. If transaction costs of one basis point are considered, none of the models generate positive returns. Thus, although the models deliver high accuracy in predicting quote updates, the predictions cannot be translated into abnormal profit due to the large impact of transaction costs on high-frequency trading. This is consistent with the efficient market hypothesis, which suggests that the reason for the predictability observed by the model is that it is costly to take these arbitrage opportunities.

However, this finding does not entirely undermine the usefulness of the model. A model with better prediction accuracy than those on offer in the existing literature may be used in two possible ways. First, a profitable trading strategy could be implemented if one could obtain relatively low transaction costs; this is more feasible for large institutional traders. Second, it could be used to help dealers abstract information from other dealers' quotes in the market. They would thereby be acting as front-runners, but in a legal and legitimate way; in anticipating other dealers' quote movements, they could make better-informed decisions regarding their own quotation strategy.

The rest of the paper is organized as follows. Section 2 briefly reviews the literature on exchange rate forecasting and introduces the ACM–ACD model in the context of high-frequency forecasting. Section 3 describes the research design and data. Section 4 presents our main results, and Section 5 concludes.

# 2 High-frequency forecasting and the ACM–ACD model

Due to data limitations, most of the early studies on exchange rate forecasting focus on lower frequency data (Dooley and Shafer, 1983; Sweeney, 1986; Levich and Thomas, 1993; Neely *et al.*, 1997). In the past two decades, advances in computer technology have improved the infrastructure of the trading platform. Due to the resultant increase in the practical importance of high-frequency trading, there is a growing literature on the predictability of price changes both at intra-day level (Osler, 2000, 2001; Curcio *et al.*, 1997; Neely and Weller, 2003; Taylor, 2007) and at ultra-high frequency in milliseconds (see Brogaard (2010) for a useful survey of these ultra-high-frequency trading rules).

One of the main findings to emerge from these studies is the importance of order-flow information to short-term exchange rate movements. Many studies document that order flow is useful in predicting exchange rate movements in both short and long horizons (e.g., Evans, 2002; Evans and Lyons, 2002; Payne, 2003). Furthermore, order flows are key media for transmitting macroeconomic announcements into exchange rate movements (Evans, 2002; Evans and Lyons, 2002). However, using order flow to predict future price presents two challenges. First, order-flow information is private to the dealers; it contains information regarding customers' expectations about future exchange rate movements. Only dealers and brokers are privileged to observe some of these order flows. Clients expect their trading interests to be treated as highly confidential and, in general, there are rules that prohibit dealers and brokers from profiting directly from the misuse of this confidential information.[7] Nevertheless, the opportunity to witness order flow can indirectly benefit dealers. It is shown that dealers do consider order flow to be informative and that they set prices accordingly (Lyons, 1995). Second, and more important, order flows are not all equally informative. Identifying which orders are the most informative requires further information regarding the identity behind the order; or, in other words, private information is required in order to interpret the observed order flow. Therefore, direct access to order-flow information without the contextual knowledge of a dealer to interpret it will not necessarily lead to more informed trading decisions.

Given that there is no publicly available data for order flow, and given the challenge of abstracting information from such data, we explore forecasting in the FX market based on dealers' quotes. In contrast to order-flow information, FX quotes – like those in other financial markets – are publicly available, since advertizing their quotes can attract order flow for dealers. We investigate whether we can abstract the private information of the dealers from the dynamic adjustments in a quote. With advances in the development of financial econometrics, we are able to explore this dynamic in two important dimensions: duration and quote changes. That is to say, we are able to jointly model when the next quote will be posted and, conditional on the duration and past quote changes, the direction in which it will change.

Quote updates by dealers are inherently in irregular time intervals. This irregularly spaced data presents a challenge to econometric modelling, since the majority of time series techniques are based on equally spaced intervals between data arrivals. As Engle and Russell (1997) point out, the choice of analysis interval is itself a very important part of the analysis, since the interval must yield a balance between noise (heteroscedasticity) and information. In this regard, Engle and Russell's (1998) autoregressive conditional duration (ACD) model is advantageous in

[7] See Harris (2003: Chapter 11) for a good discussion on legal and illegal front-runner activity in the financial trading industry.

modelling such data. They propose a point process model for inter-temporally correlated event arrival times to model data arriving at unequal time intervals.

Another characteristic found in financial data is the discreteness of price change due to the minimum tick size and the human tendency to use rounding. Russell and Engle (2005) propose a new modelling framework to jointly model price changes and durations. They decompose the joint distribution into the product of the conditional distribution of the price change and the marginal distribution of the time interval. Specifically, they model the former by means of the autoregressive conditional multinomial (ACM) model and the latter by means of the ACD model (Engle and Russell, 1998).

Following Engle and Russell (1997, 1998), we use the thinning algorithm to model the intensity of quote arrivals signifying price change. This model facilitates prediction of the length of time before prices change. Furthermore, the market microstructure model of Easley and O'Hara (1992) clearly suggests links between quote updates and the arrival of new information. Therefore, the ACM–ACD model of quote changes can also be seen as a model that captures the rate of information arrival.

# 3 Empirical design

In this section we first explain the procedure for applying the ACM–ACD model to the out-of-sample forecast. We then propose measurements and tests for evaluating forecast performance both statistically and economically. Finally, we present the exchange rate quote data used in this study.

## *3.1 Modelling and forecasting design*

We present a detailed description of the ACM–ACD model set-up in the Appendix. In this section, we focus our discussion on the application of the model to out-of-sample forecasting. In essence, our objective is to use the ACM–ACD approach to model the dynamic structure of the quote change process and use the estimated dynamics to generate forecasts. We take the following considerations into account for our empirical design.

First, we consider the options for the length of the estimation window. Zhang *et al.*, (2009) demonstrate that the ACM–ACD model captures the characteristics of the micro data very well with a sample length of 20 days, in the context of the US equity market. In comparison to the equity market, however, the FX rate market is much more liquid and fast moving. For example, there are more than 8000 quotes in the most active period of an FX trading day in our sample. Furthermore, the trading week is a natural frequency cycle in the FX market. Therefore we choose to examine daily and weekly estimation windows.

Second, we consider how often the model estimation should be updated for out-of-sample forecasting. In low-frequency contexts, modelling often uses rolling estimation, taking one observation at a time. Such a rolling method is not practical in a high-frequency context since the time per estimation is much longer than the duration between observations. Typically, for a model using one-day data it takes about five seconds to estimate the parameters, while the median duration of quote change is only about one second; thus, by the time one obtains the forecast for the theoretical next step price change, the realized price change has already moved several steps ahead. We therefore opt for a daily rolling window. The out-of-sample forecast on each day is generated using the model calibrated with data from the day before.

Specifically, at the beginning of every trading day, we use the past day's (or five days') intra-day data to fit the ACM–ACD model and obtain the model parameters. From the second quote change of the day after 2am Eastern Time, we make predictions about directional change in the next quote, by applying the model parameters with the lag duration and directional information. We record all predictions and update the model estimates at the end of the trading day for the next day's forecasts. Detail regarding the definition of a trading day is given in Section 3.3, where we discuss the data.

Finally, we consider the criteria used for recording a quoted price change. The default threshold used in our analysis is zero change; in other words, any quote update leading to a change in the mid-quote level would be recognized as an event. A threshold as low as this would encompass many events comprising very small price changes, which potentially carry less information. Furthermore, when such a model is used for forecasting, a price change signal may be correctly generated but the realized price change may yet be very small. Such forecasting may not produce economic significance. There are therefore potential gains to be had from increasing the threshold for recording events before fitting them to the models.

However, there is a danger of over-filtering, where information is significantly reduced as the result of a too-high threshold being applied. This effect is documented in Zhang *et al.* (2009), who conclude that the optimal threshold to apply to ACM–ACD forecasting in NYSE stocks is ⅛; whereas higher thresholds, such as ¼, reduce forecasting performance. We therefore explore the effect of different threshold filters on the model's performance.

However, filtering the data in out-of-sample forecasting is a precarious exercise because of the danger of data mining. Specifically, improvements in the filtering results may be due to the benefit of hindsight. We therefore need to know if, *ex ante*, a trader would have been able to learn how to filter the data over time without using full sample information. For example, Neely and Weller (2013) examine the adaptive markets hypothesis by studying the evolution of trading strategies as traders adapt their behaviour to changing circumstances. We therefore introduce a dynamic strategy to test whether switching dynamically (through learning) between the filtering strategies would improve the overall performance. To avoid look-back bias, the dynamic strategy is updated daily by comparing the past filtering-strategy returns using a rolling estimation. For each day from the 11$^{th}$ day in the sample, a hypothetical trader would compare the performance of the available filtering strategies – including different sample windows and different threshold filters – over the past five days and then choose to invest on the next day using the best of these individual strategies.

### *3.2 Evaluating forecasting performance*

We evaluate the performance of the ACM–ACD model from two angles: statistical and economic significance. First, we examine the statistical significance of the model's forecasting accuracy. To this end, we construct a variable (*Correct*) to measure the percentage of forecasts that are in the same direction as the realized quote change. Following Dacorogna *et al.* (2001), we judge whether a forecasting model is better than the random walk by examining whether the value of *Correct* is significantly higher than 50%. We approximate the variable's confidence level by the 99.99% confidence level of the random walk model.

$$\sigma_{0.01\%} \approx \frac{3.719016}{2\sqrt{n}}, \tag{1}$$

where *n* is the number of predictions. The factor 2 comes from the assumption of an equal probability of having positive or negative signals; a set-up analogous to that of evaluating the fairness of a coin by comparing coin-tossing outcomes to these confidence levels. When the value of (*Correct* minus 50%) is larger than $\sigma_{0.01\%}$, we conclude that the forecasting model is significantly better than a random walk.

In addition, we calculate the correlations between the prediction signals and the realized directions in order to evaluate the capacity of the model to predict the realized directions. We also estimate logistic models with realized direction as the dependent variable and predicted direction as the independent variable. The Wald test is used to test the significance of the coefficient on the predicted direction variable.

Tests for the predictability of models against the random walk benchmark may be biased, as Clark and West (2006) point out. In particular, they show that under the null hypothesis of no predictability, the population mean squared prediction error (MSPE) of the null 'no change' model is equal to that of the linear alternative. However, the sample MSPE of the alternative model (the model of interested) is expected to be greater than that of the null hypothesis since, under the null hypothesis of no predictability, the alternative model introduces noise into its forecasts by estimating a parameter vector that is not in fact helpful in prediction. Therefore, a finding that the null model has a smaller MSPE should not be seen as evidence against the predictability of the quote series produced by the alternative models. We therefore adopt the Clark and West testing procedure to adjust for the upward shift in the alternative model's MSPE.

Second, it is suggested that a model which has apparently performed well according to statistical evaluation criteria does not necessarily generate economic benefit (Brooks *et al.*, 2001). Our final evaluation measures the profitability of the trading strategy based on the model's forecasts. Specifically, we construct a portfolio following the buy–sell signals generated by the model.

For each day, we obtain the model parameters and apply them with the duration and direction data to create a one-step forecast. To illustrate the construction of the portfolio, we can consider the EUR/USD quote series. If the model predicts an appreciation in the EUR against the USD, we will long the EUR. Since the model parameter is estimated overnight and fixed during the day, the signals generated take very little time from receiving a quote change. We assume that one would be able to transact at the new quote midpoint. (Transaction costs are discussed later, in Section 4.2). We hold the currency until the next change of signal. The predictions are updated whenever a quote change event occurs. If a new predicted direction is different from the previous signal, we reverse our position; otherwise, we continue holding the position.

We then collect the return for each transaction. We report the average return of each transaction, the average interval of trades, and the average accumulated daily return. As a benchmark, we implement a buy-and-hold strategy whereby a currency is bought and held throughout a trading day. We show the return for holding either currency in a given pair. We calculate the returns and variance for each strategy. We compare the utility value of the portfolios by calculating the certainty-equivalent return (CER) as follows:

$$CER_{S,\gamma} = \bar{\mu}_S - \frac{\gamma}{2}\sigma_S^2 , \tag{2}$$

where $\bar{\mu}_S$ and $\sigma_S^2$ are the mean and variance of the daily returns from each strategy, *S*. $\gamma$ is the risk aversion coefficient, taking a value of 1, 3 or 5.

### *3.3 Data*

We use high-frequency bid and ask quotes from GAIN Capital[8] for spot FX rates, during the year 2010, for our three currency pairings against the dollar (USD): the euro (EUR), the British pound (GBP) and the Japanese yen (JPY).

Preliminary examination of the data shows that quotation activity is mainly concentrated between 2am and 1pm Eastern Time. As in previous studies (Neely and Weller, 2003; Rime *et al.*, 2010), we choose to model the intra-day dynamic over the period when the market is mostly widely active. This period covers the end of Asian trading hours, the whole trading session in Europe and a substantial part of US trading hours. The quotation file contains timestamps in seconds. For multiple quote changes within a second, we obtain an average mid-quote price to capture the overall information during that second.

***********************************

Table 1 about here

***********************************

Table 1 reports the summary statistics of the price and return of the exchange rate. Panel A shows the range and distribution of the exchange rate during the sample period of 2010. Panel B shows the mid-quote percentage change in basis points. N represents the number of quote changes in the sample.

Panel A of Table 1 shows that 2010 covers a wide range of exchange rate movement for the three paired currencies against the USD. We present the price level here mainly as a reference point for our later discussion of the threshold for filtering quote changes.

Panel B contains more information about the dynamics. Given the quotation style, the quote changes are calculated based on the quotes measuring the return of longing the EUR and GBP, and shorting the JPY. Overall, both the EUR and GBP depreciate against the USD during the sample period, while the JPY appreciates. The average price change between quotes is very

---

[8] GAIN capital provides an archive of historical quote data (http://ratedata.gaincapital.com/). This data contains indicative quotes, and is also featured in other studies (see, e.g., Seemann *et al.* (2011) and Chen *et al.* (2013)).

small even when expressed in basis points. The largest jumps for the EUR and GBP are in the order of 20 to 30 basis points, while quote changes at the extremes are much larger for the JPY.

# 4 Results

## 4.1 *Forecasting accuracy*

Before reporting the estimation results, we present a summary of the durations for the mid-quote movements in our sample. Table 2 summarizes the statistics under different model specifications. Durations are measured in seconds except for the maximum value, which is measured in hours. The Filter column indicates the threshold parameters used to filter the quote changes. Only movements greater than the filter parameter are recorded as quote change events. The filters in Table 2 are expressed as 10,000 times the quotation unit; for example, the threshold of 1 for the EUR implies that only when a quote change exceeds 0.0001 USD will it be recorded as a quote change event. The Est. Win. column shows the number of days used to estimate the models. The N column indicates the total number of quote change events recorded according to the filtering rules.

Table 2 shows that the mean duration is skewed by large maximum values. The median duration shows that quotes are updated typically every two to three seconds and most often every second. Filtering the price change increases durations and reduces the total number of events, as expected. When a relatively large threshold is applied, the median durations become significantly higher and the number of events drops dramatically, as shown in the last row of each currency pair. The drop in the number of events is also because we have to retain two more hours in each day to use for estimating the initial value of the day. In contrast, more modest threshold values produce small increases in duration, while the falls in the number of events are still substantial. These variations in duration and number of events provide a wide range of conditions for testing the sensitivity of the model to the information-to-noise ratio of the data.

***********************************

Tables 2 and 3 about here

***********************************

Table 3 reports the percentage of correct directional forecasts by the models. A forecasting signal is generated only when there is a change in predicted direction (that is, one requiring a change of position). For example, when a predicted quote movement changes from negative to positive, a signal of positive quote change is generated. The variable *Correct* is assigned a value of 1 when the predicted signal bears the same sign as the realized quote change, and zero otherwise.

Comparing the number of signals in this table to the number of quote changes in Table 2, we see a substantial reduction in the number of observations in Table 3. This suggests positive autocorrelation in the predictions (that is, many consecutive same-signed predictions).

The Mean column shows the mean of the variable *Correct* for all the signals generated during the sample period, which is a measure equivalent to the percentage of correct sign predictions by a given model. It shows that all the models produce a correct ratio of more than 50%, the benchmark ratio for a random walk sequence. Given the large values of N, we apply a much stricter significance level of 0.01% for testing the statistical significance of these predictions. Comparing the mean value with the confidence interval suggests that these predictions are significantly better than those of a random walk. Correlation analysis confirms that these predictions are capable of anticipating the same direction as the realized quote changes: all correlations are positive and significant, and the best model has a correlation as high as 39%. Furthermore, the Wald and Clark–West tests suggest strong predictability of the models, with all results significant at 0.01%.

These findings provide strong support for the use of the ACM–ACD approach for high-frequency exchange rate modelling. The strength of these predictions is much better than those reported in the existing literature. For example, Dacorogna *et al.* (2001: Chapter 9) report correct ratios, for their out-of-sample predictions relating to ten USD currency pairs, ranging from 50.5 to 54.2% and having a maximum value of 6% for the correlation between predictions and realized values. Furthermore, our model performance is comparable with the best results reported in Zhang *et al.* (2009), who study 50 NYSE equities in out-of-sample forecasts using the ACM–ACD model. They report a mean forecast accuracy of 54% and a maximum of 69%. They further show that prediction accuracy generally decreases as the quote update frequency increases, except for the group of stock with the highest quote update frequency, with a mean number of 780 daily updates. In comparison to equity trading, FX trading is much higher frequency and faster moving. Our findings thus complement their study, showing that the ACM–ACD model performs equally well with very high-frequency data.

When estimation windows are extended from one day to five days, the model performances are generally improved, though only marginally. This suggests that a five-day window, covering a full week's pattern of trading, contains more information than a single trading-day window. We therefore maintain a five-day estimation window for most of the models.

We find that the largest performance improvement comes from applying filters to the raw quote changes. Table 3 shows that filtering the raw data increases the information-to-noise ratio, and that there is an especially large improvement when the first step of filtering is applied. Both the mean measure of accuracy and the correlations show substantial improvement in all currency pairs. Furthermore, since the JPY's quotation unit is different, it allows us to explore more steps in the thresholds. The model performance is at its best when a filter of 100 basis point is applied in the JPY series, producing a 70% accuracy rate. However, when larger filters are applied, performances decline; indeed, those models using the largest threshold performed worse than the initial model without filtering.

As previously discussed, comparing whole-sample results for the filtering exercise can give rise to look-back bias. In practice, moreover, traders might learn from their experience and switch strategy accordingly. Thus an out-of-sample strategy would both avoid the risk of hindsight and include conditional optimization. The results of such a dynamic strategy are reported in the rows labelled 'Dynamic'. These show that the dynamic strategies do not always outperform the best filtering strategies (suggesting that these may not be achievable ex ante), but still produce improved outcomes over the non-filtered models.

Overall, the out-of-sample forecast performance of the ACM–ACD models is statistically significant and markedly superior to those in the existing literature. Forecast performance is further enhanced when modest filters are applied.

## *4.2 Economic significance*

We examine the economic significance of the forecast model by considering the possibility of profiting from these predictions.[9] We first review the before-cost portfolio return and then examine the effect of transaction costs.

### 4.2.1 Before transaction costs

Table 4 reports the performance of the long–short strategies based on the model forecasts. For reference we also record the benchmark buy-and-hold (B&H) strategy for each currency of the pairs. Table 4 shows that the mean daily returns of the prediction strategies are very high; for example, for the first three models for the EUR, the average daily accumulated returns are all higher than 10%. This is in strong contrast to the buy-and-hold strategies, where a daily return of less than 0.03% is recorded. Generally, the CERs of the strategies based on the model predictions are positive for low risk aversion, with the exception of the models that applied large thresholds; whereas all of the buy-and-hold strategies result in a negative CER. Similar contrasts in performance between our model forecasts and the buy-and-hold benchmarks are found in the other two pairs. However, we should remember that these returns do not take into account transaction costs. Before we examine the impact of transaction costs in more detail in the next

[9] The purpose of this paper is not to develop a practical trading strategy. Rather, our aim is to study price discovery in the high-frequency domain through modelling the dynamic of the midpoint. This part of our analysis explores the usefulness of the predictability based on the mid-quote modelling in the previous section. We use a simple assumption of transaction costs to examine their practical impact on the profitability of the strategy. However, transaction costs are crucial to a high-frequency trading strategy and, importantly, they vary with investor and trade characteristics. Therefore, to develop a practical high-frequency trading strategy, the detailed effect of transaction costs should be internally modelled alongside the price discovery dynamic. This is beyond the scope of the current paper. We thank the referee for highlighting this important difference in modelling strategy, which helped clarify the focus of our study.

subsection, we note that the frequency of trades is very high. A small percentage per-transaction cost would reduce the return significantly. For example, the duration of the transactions indicates that the best performance model for the EUR requires a transaction every eight seconds. This implies a requirement of about 7.5 trades per minute, 427 trades per hour or 4270 trades in the 10-hour period per day.

Interestingly, although we find that applying filters can improve forecasting accuracy, this does not necessarily translate into economic benefit. This is especially so in the case of the EUR and JPY, where the best forecasting models in terms of accuracy are those with modest thresholds, while the best return is obtained by those models without filtering. The contrasts are sharper when the CER measures are considered. Analysis suggests that strategy based on models without filtering can produce higher returns with lower standard deviations. On the other hand, models with the best forecasting accuracy may perform rather badly in economic terms. For example, the best JPY model (threshold of 100 basis point and Est. Win. of 5) in Table 3 achieves a very low return and CER in Table 4. This finding is consistent with Brooks *et al.* (2001), who study technical trading rules and find that a model that appears to have performed well according to the statistical evaluation criteria does not necessarily generate economic benefit. In the context of this research, the inferior return performance of the best directional forecast model suggests that there is asymmetry in the return for correct and incorrect sign predictions. Although the filtered models can predict direction more accurately, the size of negative impact when their prediction is wrong is much larger than the size of positive impact when they are right.

Overall, Table 4 shows that the ACM–ACD model performs well economically. Long–short strategies generate very high accumulated daily returns before transaction costs. However, the models with the best forecasting accuracy do not produce the best return distribution.

#### 4.2.2 Effect of transaction costs

Table 5 reports the effect of transaction costs on the long–short strategies based on the model forecasts. We use half a basis point as the measure of per-transaction cost in this illustration (as shown in Panel A), by taking half a basis point from each pre-cost return before accumulating the returns. We can see that this small transaction cost wipes out all profitability obtained for all the EUR and GBP models. The only models that remain positive are two of those for the JPY, using filtering. This demonstration highlights two important points. First, given the high frequency of trading required for these strategies, profitability is difficult to obtain unless very low or zero transaction costs are attainable for these algorithm trades. To our knowledge, the lowest spread in the FX market documented in public literature is one basis point in the direct inter-dealer market (Lyons, 2001). When we apply transaction costs of one basis point (as shown in Panel B), none of the above models produce a positive return. There is thus a significant barrier to entry preventing such models from being widely used in practice. This could partly explain why such high predictability and high pre-cost profitable opportunities are unarbitraged away. However, this does not negate the potential application of such forecasting models in practice. Since 2001, when Lyons wrote his book, the increased popularity of electronic platforms (BIS, 2013) may have driven transaction costs down to a lower level. More importantly, a model with such high levels of accuracy in predicting quote changes would be useful to dealers, who are obligated to make the market with their high-frequency quote updates. For example, if a downtick forecast is generated by the model for the USD/EUR pair, a dealer who has a net position in the EUR would want to reduce her position by posting a lower ask (selling) price but not as low as the expected price move. In this way, she would be able to attract order flow from customers buying the EUR and selling at a better price.

Second, the results in Table 5 again show the advantage of data filtering. In the sections above, we found that filtering raw data with modest thresholds improves forecasting accuracy. However, it does not necessarily generate high pre-cost profit. When transaction costs are considered, trading frequency plays a much more important role, given its compounding effect. It shows that a second benefit of applying a threshold, in addition to the increase in the information-to-noise ratio in the point process, is to reduce the number of transactions and hence lower the transaction costs of the trading strategy.

Overall, we show that transaction costs have a detrimental effect on the profitability of high-frequency trading strategies. On the one hand, this finding supports market efficiency, in the sense that no abnormal profit can be systematically generated using a prediction model. The existence of predictability is partly driven by the transaction cost, which prevents information being fully reflected in the price.[10] On the other hand, the finding highlights an important condition for the practical application of these prediction models: if they are to be used for trading strategy, profitability is possible only with very low transaction costs.

It is important to note that the above discussion is from the perspective of traders, who have to *pay* the transaction costs. The return before costs suggests a great potential economic benefit to dealers, in contrast, who make the market and *earn* the transaction costs. When applying this model to the prevailing quote data in the market, the dealer can anticipate market movement and provide a more competitive quote to attract greater order flow.

## 5 Conclusion

This paper takes up the challenge of forecasting exchange rate movement tick by tick. We use the ACM–ACD approach (Russell and Engle, 2005), a high-frequency financial econometric model capable of measuring a marked point process with irregular arrival times, to examine the accuracy and economic significance of the resulting forecasts in the context of the FX rate market. Specifically, we study the quote changes of the three most traded currency pairs: the EUR, JPY and GBP against the USD, during the one-year period of 2010. Our empirical analyses produce the following findings.

First, the ACM–ACD model provides much improved forecasting over the benchmark random walk model and those models reported in the existing literature. The forecast accuracies of the models studied range from 54 to 70%. Correlation analyses, the Wald test and the Clark–

[10] The previous literature mainly argues that transaction costs cause the slow adjustment to a small mispricing (e.g., Sercu *et al.*, 1995; Panos *et al.*, 1997; Roll *et al.*, 2007 and Oehmke, 2011).

West test on the relationship between predictions and realized quote movements also confirm the positive and significant predictive power of the models against the random walk benchmark.

Second, we find that a trading strategy based on the directional forecasts could earn substantial returns before transaction costs. However, when transaction costs of half a basis point are applied to each intra-day transaction, the profits quickly turn into loss for most of the model predictions except for two filtered models for the. If transaction costs of one basis point are considered, none of the models generate positive returns. Thus, although these models produce high accuracy in predicting quote updates, the predictions cannot turn into abnormal profit due to the high impact of transaction costs on high-frequency trading. This is consistent with the efficient market hypothesis; the reason why these prediction patterns exist is that they cannot be arbitraged away due to cost.

However, these findings do not preclude the future application of the ACM–ACD model in practice. The increased use of electronic trading platforms has in general driven transaction costs down (BIS, 2013). Hence, there is the possibility that some institutions, which can obtain transaction costs as low as half a basis point, will be able to enjoy some economic benefit from this forecast model. Furthermore, the model could certainly provide advantageous information regarding price movement to FX dealers, when applied to a series of quote changes that contain multiple dealers' quotes. Dealers are obligated to quote on both sides of the market in order to attract order flow. A good forecast model would assist them in both maximizing the gains from the spread and managing their inventory risk by setting quotes that are of benefit to them, given the forecast price movement in the next tick.

Third, we examine the impact of data filtering on the model's forecasting performance. The ACM–ACD model and especially the ACD model are designed to model an event that arrives in irregular time intervals. The realized sequence such as the quote changes may contain much noise; for example, there may be quote updates that contain no changes in the quote price. Engle and Russell (1997) point out the importance of filtering the data in empirical applications of the ACD model. We take this further and examine the impact of higher thresholds on the marked point process. When applying a threshold to the quote change, it effectively redefines the events, requiring that only quote changes larger than the threshold parameters be recorded as event points; the marked point process is constructed accordingly.

Our empirical study shows two important aspects of data filtering with a threshold parameter. First, applying a modest threshold will increase the information-to-noise ratio, while too high a threshold will reduce the information content of the point process. There is thus a balance between the reduction of noise and the reduction of information when performing data filtering; seeking the optimal threshold is an empirical challenge in itself. Second, from the perspective of trading strategy, there is an important benefit in using thresholds, in addition to the improved prediction accuracy: it reduces the number of transactions required, given the longer durations between the signals generated by the model. A smaller number of transactions means lower transaction costs overall; an especially significant consideration given the compounding effect of per-trade transaction costs on the overall profitability of the strategy.

Overall, we show that the ACM–ACD approach produces good out-of-sample forecasts when applied to the financial market with highest-frequency trading: the FX rate market. Its advanced dynamic modelling structure appears capable of abstracting private information from the publicly observable data. However, it is only possible to benefit economically from these forecasts when very low transaction costs are obtainable. These findings support the efficient market hypothesis in general, while providing empirical evidence that reveals the potential of the practical applications of the model for large institutional FX traders and dealers.

## Appendix: The ACM-ACD model set-up

The detail of the ACM–ACD model settings is set out below. Following Russell and Engle's (2005) notation, let $t_i$ denote the time that the $t^{th}$ market event occurs. In the context of the current research, a market event is a change of the mid-quote price. Let $\tau_i$ denote the duration between two market events occurring at times $t_{i-1}$ and $t_i$. And let $y_i$ denote the price change of this market event.

The choices for construction of the $y_i$ and $\tau_i$ series are various. $y_i$ can have two dimensions, comprising only up and down tick price movements; three dimensions, comprising up, down and zero tick movements; or five dimensions, as in Russell and Engle (2005) comprising up two ticks, up one tick, zero ticks, down one tick and down two ticks. $\tau_i$ can measure trade duration or quote revision duration. Furthermore, construction of the price movement $y_i$ and duration $\tau_i$ series can be conditional on the thresholds of other variables. The most obvious choice of threshold variable is the price change itself. For example, $y_i$ only contains those points with price movements larger than a given number of ticks.

Zhang *et al.* (2009) point out two advantages of choosing quote change as a market event. First, use of the mid-quote change removes the microstructure effect of bid–ask bounce and therefore reduces the noise in the data. Second, forecasting mid-quote change has a practical advantage over forecasting trade price change, in that the forecast midpoint constitutes a useful reference point for investors, regardless of whether they are buyers or sellers.

In addition, our discussion in Section 2 points out that a potential benefit of modelling the dealers' quote change behaviour is that information private to the dealers may be abstracted from duration modelling. Therefore, we choose to model the mid-quote change and duration in our study.

We let $y_i$ denote the direction of change of mid-quote price, giving $y_i$ two dimensions[11]:

$y_i = 1$ if the mid-quote price increases over duration $\tau_i$ ,and

$y_i = -1$ if the mid-quote price decreases over duration $\tau_i$ .

The ACM–ACD model decomposes the joint conditional density of $y_i$ and $\tau_i$ into the product of the conditional density of the mark and the marginal density of the arrival times, both conditional on the past information set:

$$f(y_i, \tau_i \mid y^{(i-1)}, \tau^{(i-1)}) = g(y_i \mid y^{(i-1)}, \tau^{(i)}) q(\tau_i \mid y^{(i-1)}, \tau^{(i-1)}), \qquad \text{(A1)}$$

where $g(\cdot)$ denotes the probability density function of the price change $y_i$ conditional on $\tau^{(i)}$ and $y^{(i-1)}$, and $q(\cdot)$ denotes the density function of the $i^{th}$ quote duration conditional on $\tau^{(i-1)}$ and $y^{(i-1)}$.

Russell and Engle (2005) use the autoregressive conditional multinomial (ACM) specification to model $g(\cdot)$. In the two-dimension case, let $\tilde{x}_i$ be a 2*1 vector indicating the possible states of the discrete mid-quote price change $y_i$. $\tilde{x}_i$ takes the $j^{th}$ column of the 2*2 identity matrix if the $j^{th}$ state occurred. Let $\tilde{\pi}_i$ denote a 2*1 vector of conditional probabilities associated with the states. The $j^{th}$ element of $\tilde{\pi}_i$ corresponds to the probability that the $j^{th}$ element of $\tilde{x}_i$ takes the value 1. $\tilde{\pi}_i$ must satisfy the following two conditions:

1. All elements of $\tilde{\pi}_i$ are non-negative.
2. All columns must sum to unity.

Russell and Engle (2005) propose using an inverse logistic transformation that imposes such conditions directly. Let $\pi_i$ denote the probability of an up movement of mid-quote price, and define the log of the probability ratios as $h(\pi_i) = \ln(\pi_i / (1 - \pi_i))$. The ACM model assumes that $h(\pi_i)$ follows an autoregressive conditional specification:

[11] For simplicity, in the context of our model we use 'mid-quote price change' and 'price change' interchangeably.

$$h(\pi_i) = c + \sum_{j=1}^{p} a_j (x_{i-j} - \pi_{i-j}) + \sum_{j=1}^{q} b_j h(\pi_{i-j}) + \chi_1 \ln(\tau_i) + \chi_2 \ln(\tau_{i-1}) . \tag{A2}$$

The conditional probability of an up movement of mid-quote price is easily recovered from the above inverse logistic transformation using the following equation:

$$\pi_i = \frac{\exp[c + \sum_{j=1}^{p} a_j (x_{i-j} - \pi_{i-j}) + \sum_{j=1}^{q} b_j h(\pi_{i-j}) + \chi_1 \ln(\tau_i) + \chi_2 \ln(\tau_{i-1})]}{1 + \exp[c + \sum_{j=1}^{p} a_j (x_{i-j} - \pi_{i-j}) + \sum_{j=1}^{q} b_j h(\pi_{i-j}) + \chi_1 \ln(\tau_i) + \chi_2 \ln(\tau_{i-1})]} . \tag{A3}$$

And the probability of a down movement of mid-quote price is calculated by $1 - \pi_i$. Given the initial condition, the entire path of $\tilde{\pi}_i$ can be constructed, and the log likelihood is expressed as:

$$L = \sum_{i=1}^{N} \sum_{j=1}^{K} (\tilde{x}_{ij} \ln(\tilde{\pi}_{ij})) = \sum_{i=1}^{N} \tilde{x}_i' \ln(\tilde{\pi}_i) . \tag{A4}$$

In the above equation, $\tilde{x}_i$ represents the realized observations and $\tilde{\pi}_i$ contains the variables needing to be estimated. In Russell and Engle (2005), durational dynamics are specified by an exponential ACD model. This means that the hazard function is constant. Here, the conditional density function is generalized to the Weibull Log–ACD model. The log likelihood associated with the $i^{th}$ duration is:

$$\ln(q(\tau_i \mid y^{(i-1)}, \tau^{(i-1)})) = \ln(\frac{\gamma}{\tau_i}) + \gamma \ln(\frac{\Gamma(1+1/\gamma)\tau_i)}{\psi_i}) - (\frac{\Gamma(1+1/\gamma)\tau_i)}{\psi_i})^{\gamma} , \tag{A5}$$

where:

$$\ln(\psi_i) = \omega + \sum_{j=1}^{u} \alpha_j \varepsilon_{i-j} + \sum_{j=1}^{v} \beta_j \ln(\psi_{i-j}) . \tag{A6}$$

The estimation of the ACM–ACD model can be performed by maximizing the sum of the ACM and ACD log-likelihood functions or separately maximizing the two log-likelihood functions.

## References


Bacchetta, P. and van Wincoop, E., 2004, 'A scapegoat model of exchange-rate fluctuations', *American Economics Review*, Vol. 94, pp. 114–18.

Bauwens, L. and Giot, P., 2003, 'Asymmetric ACD models: introducing price information in ACD models', *Empirical Economics*, Vol. 28, 709–31.

Brogaard, J.A., 2010, 'High frequency trading and its impact on market quality', Working paper (Northwestern University).

BIS, 2013, 'Triennial Central Bank Survey of foreign exchange and derivatives market activity in 2013' Bank for International Settlements, Date last accessed: May 2014, https://www.bis.org/publ/rpfx13.htm

Brooks, C., Rew, A.G. and Ritson, S., 2001, 'A trading strategy based on the lead-lag relationship between the spot index and futures contract for the FTSE100', *International Journal of Forecasting*, Vol. 17, pp. 31–4.

Chen, S., Chou, C. and Chen, N., 2013, 'A wavelet transform analysis of the relationship between unexpected macroeconomic news and foreign exchange rates', *Applied Economics Letters*, Vol. 20, pp. 292–6.

Cheung, Y.W., Chinn, M.D. and Garcia Pascual, A., 2005, 'Empirical exchange rate models of the nineties: are any fit to survive?', *Journal of International Money and Finance*, Vol. 24, pp. 1150–75.

Clark T.E. and West K.D., 2006, 'Using out-of-sample mean squared prediction errors to test themartingale difference hypothesis', *Journal of Econometrics*, Vol. 135, pp. 155–86.

Curcio, R., Goodhart, C., Guillaume, D. and Payne, R., 1997, 'Do technical trading rules generate profits? Conclusions from the intra-day foreign exchange market', *International Journal of Finance and Economics*, Vol. 2, pp. 267–80.

Dacorogna, M.M., Gencay, R., Muller, U., Olsen, R.B. and Picktet, O.V., 2001, *An introduction to high frequency finance* (Academic Press, New York).

Dionne, G., Duchesne, P. and Pacurar, M., 2009, 'Intraday Value at Risk (IVaR) using tick-by-tick data with application to the Toronto Stock Exchange', *Journal of Empirical Finance*, Vol. 16, pp. 777–92.

Dooley, M.P. and Shafer, J.R., 1983, 'Analysis of short-run exchange rate behavior: March 1973 to November 1981', in T. Bigman and T. Taya, eds, *Exchange Rate and Trade Instability: Causes, Consequences and Remedies* (Ballinger Publishing Company, Pensacola), pp. 43–69.

Easley, D. and O'Hara, M., 1987, 'Price, trade size, and information in securities markets', *Journal of Financial Economics*, Vol. 19, pp. 69–90.

Easley, D. and O'Hara M., 1992, 'Time and the process of security price adjustment', *Journal of Finance*, Vol. 47, pp. 577–605.

Engel, C., Mark, N.C. and West, K.D., 2008, 'Exchange rate models are not as bad as you think', in D. Acemoglu, K. Rogoff, and M. Woodford, eds, *NBER Macroeconomics Annual 2007* (MIT Press, Cambridge MA), pp. 381–441.

Engle, R. and Russell, J., 1997, 'Forecasting the frequency of changes in quoted foreign exchange prices with the autoregressive conditional duration model', *Journal of Empirical Finance*, Vol. 12, pp. 187–212.

Engle, R. and Russell, J., 1998, 'Autoregressive conditional duration: a new approach for irregularly spaced transaction data', *Econometrica*, Vol. 66, pp. 1127–62.

Evans, M.D.D., 2002, 'FX trading and exchange rate dynamics', *Journal of Finance*, Vol. 57, pp. 2405–47.

Evans, M.D.D. and Lyons, R.K., 2002, 'Order flow and exchange rate dynamics', *Journal of Political Economics*, Vol. 110, pp. 170–80.

Evans, M.D.D. and Lyons, R.K., 2005, 'Meese–Rogoff redux: micro-based exchange-rate forecasting', *American Economics Review*, Vol. 95, pp. 405–14.

Evans, M.D.D. and Lyons, R.K., 2007, 'Exchange rate fundamentals and order flow', Working paper No. 13151 (NBER).

Harris, L., 2003, *Trading and exchanges: market microstructure for practitioners* (Oxford University Press, New York).

Hong, Y., Li, H. and Zhao, F., 2007, 'Can the random walk model be beaten in out-of-sample density forecasts? Evidence from intraday foreign exchange rates', *Journal of Econometrics*, Vol. 141, pp. 736–76.

Levich, R.M. and Thomas, L.R., 1993, 'The significance of technical trading rule profits in the foreign exchange market: a bootstrap approach', *Journal of International Money and Finance*, Vol. 12, pp. 451–74.

Lyons, R.K., 1995, 'Tests of microstructural hypotheses in the foreign exchange market', *Journal of Financial Economics*, Vol. 39, pp. 321–51.

Lyons, R.K., 2001, *The Microstructure Approach to Exchange Rates* (MIT Press, Cambridge MA).

Meese, R.A. and Rogoff, K., 1983, 'Empirical exchange rate models of the seventies: Do they fit out of sample?', *Journal of International Economics*, Vol. 14, pp. 3–24.

Neely, C.J. and Weller, P.A., 2003, 'Intraday technical trading in the foreign exchange market', *Journal of International Money and Finance*, Vol. 22, pp. 223–37.

Neely, C.J. and Weller, P.A., 2013, 'Lessons from the evolution of foreign exchange trading strategies', *Journal of Banking and Finance*, Vol. 37, pp. 3783–98.

Neely, C.J. and Weller, P.A., Dittmar, R., 1997, 'Is technical analysis in the foreign exchange market profitable? A genetic programming approach', *Journal of Financial and Quantitative Analysis*, Vol. 32, pp. 405–26.

Oehmke, M., 2011, 'Gradual arbitrage', Working paper, Columbia University.

Osler, C.L., 2000, 'Support for resistance: technical analysis and intraday exchange rates', *Federal Reserve Bank of New York Policy Review*, Vol. 6, pp. 253–68.

Osler, C.L., 2001, 'Conditional orders and exchange-rate dynamics: explaining the success of technical analysis', Unpublished manuscript (Federal Reserve Bank of New York).

Panos, A.M., Nobay, R. and Peel, D.A., 1997, 'Transactions costs and nonlinear adjustment in real exchange rates: an empirical investigation', *Journal of Political Economy*, Vol. 105, pp. 862–79.

Payne, R., 2003, 'Informed trade in spot foreign exchange markets: an empirical investigation', *Journal of International Economics*, Vol. 61, pp. 307–29.

Rime, D., Sarno, L. and Sojli, E., 2010, 'Exchange rate forecasting, order flow and Macroeconomic information', *Journal of International Economics*, Vol. 80, pp. 72–88.

Roll, R., Schwartz, E. and Subrahmanyam, A., 2007, 'Liquidity and the law of one price: the case of the futures-cash basis', *Journal of Finance*, Vol. 62, pp. 2201–34.

Russell, J., and Engle, R., 2005, 'A discrete-state continous-time model of financial transactions prices and times: the autoregressive conditional multinomial–utoregressive conditional duration model', *Journal of Business and Economic Statistics*, Vol. 23, pp. 166–80.

Sarno, L. and Taylor, M.P., 2002, 'Purchasing power parity and the real exchange rate,' *IMF Staff Papers*, Vol.49, 65–105.

Sarno, L. and Valente, G., 2005, 'Empirical exchange rate models and currency risk: some evidence from density forecasts', *Journal of International Money and Finance*, Vol. 24, pp. 363–85.

Seemann, L., McCauley, J.L. and Gunaratne, G.H., 2011, 'Intraday volatility and scaling in high frequency foreign exchange markets', *International Review of Financial Analysis*, Vol. 20, pp. 121–6.

Sercu, P., Uppal, R. and Hulle, C.V., 1995, 'The exchange rate in the presence of transaction costs: implications for tests of purchasing power parity', *Journal of Finance*, Vol. 50, pp. 1309–19.

Sweeney, R.J., 1986, 'Beating the foreign exchange market', *Journal of Finance*, Vol. 41, pp. 163–82.

Taylor, A.M. and Taylor, M.P., 2004, 'The purchasing power parity debate', *Journal of Economic Perspectives*, Vol. 18, pp. 135–58.

Taylor, N., 2007, 'A new econometric model of index arbitrage', *European Financial Management,* Vol. 13, pp. 159-183.

Zhang, Q., Cai, C. and Keasey, K., 2009, 'Forecasting using high-frequency data: a comparison of asymmetric financial duration models', *Journal of Forecasting*, Vol. 28, pp. 371–86.

**Table 1. Summary statistics of exchange rate price and return.**

This table reports the summary statistics of the price and return of the exchange rate. Panel A reports the range and distribution of the exchange rate during the sample period of 2010. Panel B reports the mid-quote percentage change in basis points. N represents the number of quote changes in the sample.

| **Currency** | **Mean** | **Min** | **P25** | **Median** | **P75** | **Max** | **SD** | **N** |
|---|---|---|---|---|---|---|---|---|
| Panel A. Price (mid-quote) | | | | | | | | |
| EUR | 1.32355 | 1.19025 | 1.27540 | 1.33025 | 1.36725 | 1.45800 | 0.06027 | 2,052,284 |
| GBP | 1.53786 | 1.42313 | 1.50108 | 1.54100 | 1.57830 | 1.64580 | 0.05145 | 2,890,427 |
| JPY | 88.0628 | 80.3350 | 84.3500 | 88.7850 | 91.3350 | 94.9900 | 3.89761 | 1,306,350 |
| Panel B. Return (percentage change in basis points) | | | | | | | | |
| EUR | −0.00031 | −22.24019 | −0.74566 | −0.34743 | 0.74685 | 20.34505 | 0.77914 | 2,052,025 |
| GBP | −0.00005 | −26.63635 | −0.48042 | 0.00000 | 0.48099 | 29.00512 | 0.57611 | 2,890,168 |
| JPY | −0.00110 | −102.33057 | −1.07708 | 0.00000 | 1.07724 | 53.32961 | 1.01469 | 1,306,091 |

**Table 2. Summary of realized durations.**

This table presents a summary of durations for the mid-quote movements in our sample, shown under the different model specifications. The Filter column indicates the threshold parameters used to filter the quote changes; only movements greater than the filter parameter are recorded as quote change events. The Est. Win. column reports the number of days used in the model estimation. All durational statistics are measured in seconds except for the maximum value which is in hours. N represents the total number of movement events recorded. The sample period is the calendar year 2010.

| Currency | Filter | Est. Win. | Mean | Min | Median | Max (hour) | Mode | N |
|---|---|---|---|---|---|---|---|---|
| EUR | 0 | 1 | 5 | 1 | 3 | 0.25 | 1 | 2,052,025 |
| | 0 | 5 | 5 | 1 | 3 | 0.25 | 1 | 2,025,474 |
| | 1 | 5 | 6 | 1 | 3 | 0.27 | 1 | 1,681,113 |
| | 1.5 | 5 | 73 | 1 | 26 | 0.96 | 1 | 110,650 |
| GBP | 0 | 1 | 4 | 1 | 2 | 0.16 | 1 | 2,890,168 |
| | 0 | 5 | 4 | 1 | 2 | 0.16 | 1 | 2,840,714 |
| | 1 | 5 | 7 | 1 | 4 | 0.33 | 1 | 1,417,434 |
| | 1.5 | 5 | 60 | 1 | 22 | 1.72 | 1 | 135,122 |
| JPY | 0 | 1 | 8 | 1 | 3 | 0.33 | 1 | 1,306,091 |
| | 0 | 5 | 8 | 1 | 3 | 0.33 | 1 | 1,286,957 |
| | 75 | 5 | 12 | 1 | 5 | 0.33 | 1 | 832,939 |
| | 100 | 5 | 22 | 1 | 7 | 0.94 | 1 | 364,772 |
| | 150 | 5 | 152 | 1 | 37 | 3.68 | 1 | 43,332 |

**Table 3. Direction quality and signal correlations.**

This table reports the percentage of correct directional forecasts by the models, as measured by the variable *Correct*. This variable is assigned a value of 1 when the predicted sign of quote change is the same as that of the realized quote change, and zero otherwise. A prediction signal is generated only when there is a change of predicted direction. The Mean column reports the mean of the variable *Correct* for all the signals generated during the sample period. The two Confidence interval columns report the approximation of the confidence interval of the random walk at 0.01% and 99.99%, respectively, which is calculated using Equation (1). The Correlation column reports the Pearson correlation. The Wald ChiSq column reports the statistics for the Wald test on the predictability of the forecast against the realized quote changes. The Clark–West column reports the test evaluating the null that a given series follows a zero-mean Martingale difference, using out-of-sample MSPE as proposed by Clark and West (2006). All model specifications using different filter thresholds and estimation windows are reported. *** signifies statistical significance at 0.01%. N represents the number of signals generated by each model.

| Currency | Filter | Est. Win. | Mean | Confidence interval | | Correlation | Wald ChiSq | Clark–West | N |
|---|---|---|---|---|---|---|---|---|---|
| | | | | 0.0001 | 0.9999 | | | | |
| EUR | 0 | 1 | 0.599 | 0.498 | 0.502 | 0.20 *** | *** | *** | 1,123,006 |
| | 0 | 5 | 0.603 | 0.498 | 0.502 | 0.21 *** | *** | *** | 1,197,581 |
| | 1 | 5 | 0.617 | 0.498 | 0.502 | 0.23 *** | *** | *** | 987,905 |
| | 1.5 | 5 | 0.538 | 0.489 | 0.511 | 0.08 *** | *** | *** | 29,197 |
| | Dynamic | | 0.614 | 0.498 | 0.502 | 0.19 *** | *** | *** | 977,844 |
| GBP | 0 | 1 | 0.553 | 0.499 | 0.501 | 0.11 *** | *** | *** | 1,549,743 |
| | 0 | 5 | 0.560 | 0.498 | 0.502 | 0.12 *** | *** | *** | 1,063,140 |
| | 1 | 5 | 0.589 | 0.498 | 0.502 | 0.18 *** | *** | *** | 693,357 |
| | 1.5 | 5 | 0.541 | 0.491 | 0.509 | 0.08 *** | *** | *** | 44,588 |
| | Dynamic | | 0.599 | 0.498 | 0.502 | 0.16 *** | *** | *** | 630,161 |
| JPY | 0 | 1 | 0.629 | 0.498 | 0.502 | 0.26 *** | *** | *** | 750,633 |
| | 0 | 5 | 0.634 | 0.498 | 0.502 | 0.27 *** | *** | *** | 804,943 |
| | 75 | 5 | 0.676 | 0.497 | 0.503 | 0.35 *** | *** | *** | 549,199 |
| | 100 | 5 | 0.695 | 0.496 | 0.504 | 0.39 *** | *** | *** | 240,622 |
| | 150 | 5 | 0.529 | 0.484 | 0.516 | 0.06 *** | *** | *** | 13,542 |
| | Dynamic | | 0.663 | 0.497 | 0.503 | 0.29 *** | *** | *** | 500,055 |

**Table 4. Performance of long–short strategies based on model forecasts.**

This table reports the daily accumulated return, certainty equivalent return (CER) and duration of transactions. Returns are reported in basis points. CER is calculated using Equation (2) with the risk aversion parameter $\gamma$ set to 1, 3 and 5 respectively. Duration is measured in seconds (except for the buy and hold strategies which is in hours). All model specifications using different filter thresholds and estimation windows are reported. In addition, the returns of buy-and-hold (B&H) strategies for each currency pair are reported. N represents the number of days.

| Currency | Filter | Est. Win. | Return | | | | | CER | | | Duration | | N |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | Mean | Min | Median | Max | SD | (1) | (3) | (5) | Mean | Median | |
| EUR | 0 | 1 | 1107.2 | −750 | 1201 | 2220 | 484 | 865 | 381 | −104 | 9 | 5 | 259 |
| | 0 | 5 | 1251.6 | −133 | 1251 | 2220 | 326 | 1089 | 763 | 437 | 8 | 5 | 255 |
| | 1 | 5 | 1001.6 | −473 | 1030 | 1831 | 345 | 829 | 484 | 139 | 10 | 6 | 255 |
| | 1.5 | 5 | 9.7 | −133 | 0 | 321 | 61 | −21 | −82 | −144 | 230 | 72 | 255 |
| | Dynamic | | 1019.6 | −473 | 1058 | 1857 | 387 | 826 | 439 | 51 | 10 | 6 | 250 |
| | B&H: EUR | | −2.6 | −161 | −2 | 178 | 55 | −30 | −85 | −139 | 10 hr | | 259 |
| | B&H: USD | | 2.9 | −175 | 2 | 164 | 55 | −24 | −79 | −134 | 10 hr | | 259 |
| GBP | 0 | 1 | 561.9 | −1562 | 516 | 1983 | 541 | 291 | −250 | −792 | 7 | 4 | 259 |
| | 0 | 5 | 465.2 | −575 | 248 | 1637 | 517 | 207 | −311 | −828 | 8 | 4 | 255 |
| | 1 | 5 | 515.9 | −158 | 476 | 1514 | 393 | 320 | −73 | −466 | 14 | 7 | 255 |
| | 1.5 | 5 | 37.8 | −100 | 25 | 598 | 83 | −4 | −87 | −171 | 166 | 61 | 255 |
| | Dynamic | | 481.3 | −100 | 446 | 1514 | 428 | 267 | −160 | −588 | 14 | 5 | 250 |
| | B&H: GBP | | −0.6 | −141 | −2 | 122 | 52 | −27 | −79 | −132 | 10 hr | | 259 |
| | B&H: USD | | 0.8 | −121 | 2 | 143 | 52 | −25 | −78 | −130 | 10 hr | | 259 |
| JPY | 0 | 1 | 1314.2 | −580 | 1453 | 2663 | 693 | 968 | 274 | −419 | 14 | 6 | 259 |
| | 0 | 5 | 1489.3 | −78 | 1600 | 2693 | 572 | 1203 | 631 | 59 | 13 | 6 | 255 |
| | 75 | 5 | 1265.0 | −384 | 1459 | 2507 | 702 | 914 | 212 | −490 | 18 | 8 | 255 |
| | 100 | 5 | 535.1 | −529 | 574 | 1407 | 377 | 347 | −30 | −408 | 33 | 10 | 255 |
| | 150 | 5 | −1.0 | −162 | 0 | 530 | 64 | −33 | −96 | −160 | 378 | 80 | 255 |
| | Dynamic | | 923.1 | −561 | 801 | 2507 | 541 | 653 | 112 | −429 | 18 | 7 | 250 |
| | B&H: USD | | −5.7 | −216 | −6 | 157 | 52 | −31 | −83 | −135 | 10 hr | | 259 |
| | B&H: JPY | | 5.9 | −155 | 6 | 221 | 52 | −20 | −72 | −124 | 10 hr | | 259 |

**Table 5. Effect of transaction costs.**

This table reports the effect of transaction costs on the long–short strategies based on the model forecasts. Panel A shows the impact when half a basis point is taken away from each pre-cost return before accumulating them. Panel B shows the corresponding impact of transaction costs of one basis point. The daily return, CER and the duration of transactions are reported.

| Currency | Filter | Est. Win. | Return | | | | | CER | | | N |
|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | Mean | Min | Median | Max | SD | (1) | (3) | (5) | |
| Panel A. Transaction costs: half basis point | | | | | | | | | | | |
| EUR | 0 | 1 | −1050 | −3168 | −975 | 59 | 531 | −1315 | −1846 | −2377 | 259 |
| | 0 | 5 | −1094 | −3168 | −1029 | −165 | 461 | −1325 | −1786 | −2247 | 255 |
| | 1 | 5 | −929 | −2800 | −868 | 31 | 438 | −1148 | −1586 | −2023 | 255 |
| | 1.5 | 5 | −47 | −342 | −44 | 132 | 66 | −80 | −145 | −211 | 255 |
| | Dynamic | | −930 | −2800 | −870 | 67 | 475 | −1168 | −1643 | −2118 | 250 |
| GBP | 0 | 1 | −2119 | −5282 | −1876 | −20 | 955 | −2596 | −3551 | −4506 | 259 |
| | 0 | 5 | −1440 | −4356 | −1359 | 66 | 1106 | −1993 | −3099 | −4205 | 255 |
| | 1 | 5 | −817 | −1989 | −826 | 49 | 425 | −1029 | −1454 | −1878 | 255 |
| | 1.5 | 5 | −49 | −285 | −38 | 252 | 76 | −87 | −163 | −239 | 255 |
| | Dynamic | | −740 | −3463 | −800 | 252 | 693 | −1087 | −1780 | −2473 | 250 |
| JPY | 0 | 1 | −209 | −3363 | −1 | 713 | 633 | −525 | −1159 | −1792 | 259 |
| | 0 | 5 | −177 | −3364 | 41 | 713 | 655 | −505 | −1160 | −1815 | 255 |
| | 75 | 5 | 105 | −2569 | 159 | 941 | 432 | −111 | −543 | −974 | 255 |
| | 100 | 5 | 46 | −1727 | 39 | 464 | 225 | −66 | −292 | −517 | 255 |
| | 150 | 5 | −27 | −490 | −13 | 148 | 63 | −59 | −122 | −186 | 255 |
| | Dynamic | | −109 | −1820 | 76 | 941 | 598 | −408 | −1006 | −1604 | 250 |
| Panel B. Transaction costs: one basis point | | | | | | | | | | | |
| EUR | 0 | 1 | −2746 | −5721 | −2739 | 50 | 953 | −3222 | −4175 | −5127 | 259 |
| | 0 | 5 | −2931 | −5721 | −2853 | −1048 | 699 | −3281 | −3980 | −4680 | 255 |
| | 1 | 5 | −2502 | −4564 | −2513 | 16 | 706 | −2855 | −3562 | −4268 | 255 |
| | 1.5 | 5 | −103 | −593 | −86 | 67 | 102 | −154 | −256 | −358 | 255 |
| | Dynamic | | −2507 | −5173 | −2549 | 0 | 829 | −2921 | −3750 | −4579 | 250 |
| GBP | 0 | 1 | −4071 | −7362 | −3937 | −24 | 1221 | −4681 | −5902 | −7123 | 259 |
| | 0 | 5 | −2863 | −6620 | −3156 | 26 | 1892 | −3809 | −5701 | −7593 | 255 |
| | 1 | 5 | −1938 | −3874 | −2042 | 35 | 940 | −2408 | −3348 | −4288 | 255 |
| | 1.5 | 5 | −135 | −489 | −117 | 116 | 115 | −192 | −307 | −421 | 255 |
| | Dynamic | | −1731 | −5746 | −2154 | 44 | 1437 | −2449 | −3886 | −5322 | 250 |
| JPY | 0 | 1 | −1499 | −5562 | −1309 | 0 | 870 | −1935 | −2805 | −3675 | 259 |
| | 0 | 5 | −1582 | −5562 | −1369 | −116 | 865 | −2015 | −2880 | −3745 | 255 |
| | 75 | 5 | −917 | −4258 | −833 | −22 | 538 | −1186 | −1724 | −2262 | 255 |
| | 100 | 5 | −411 | −2773 | −347 | 65 | 364 | −594 | −958 | −1323 | 255 |
| | 150 | 5 | −53 | −827 | −25 | 118 | 91 | −98 | −189 | −281 | 255 |
| | Dynamic | | −1017 | −3677 | −642 | 65 | 910 | −1472 | −2381 | 3291 | 250 |